Lensed CMB simulation and parameter estimation

Antony Lewis[1,*]

[1]CITA, 60 St. George St, Toronto M5S 3H8, ON, Canada.

Modelling of the weak lensing of the CMB will be crucial to obtain correct cosmological parameter constraints from forthcoming precision CMB anisotropy observations. The lensing affects the power spectrum as well as inducing non-Gaussianities. We discuss the simulation of full sky CMB maps in the weak lensing approximation and describe a fast numerical code. The series expansion in the deflection angle cannot be used to simulate accurate CMB maps, so a pixel remapping must be used. For parameter estimation accounting for the change in the power spectrum but assuming Gaussianity is sufficient to obtain accurate results up to Planck sensitivity using current tools. A fuller analysis may be required to obtain accurate error estimates and for more sensitive observations. We demonstrate a simple full sky simulation and subsequent parameter estimation at Planck-like sensitivity. The lensed CMB simulation and parameter estimation codes are publicly available.

arXiv:astro-ph/0502469v2 12 Apr 2005

I. INTRODUCTION

The CMB temperature and polarization anisotropies are being measured with ever more precision. The statistics of the anisotropies can provide valuable limits on cosmological parameters as well as constrain early universe physics. As we enter the era of precision measurement with signal-dominated observations out to small scales, the non-linear effects will become important. One of the most important of these on scales of most interest for parameter estimation is that of weak lensing. Fortunately it can be modelled accurately using linear physics: the linear potentials along the line sight lensing the linear perturbations at the last scattering surface [1, 2, 3]. Modelling of fully non-linear evolution is not required for the near future on scales with $l \lesssim 2000$, and non-linear corrections can be applied to the lensing potential if and when required. On smaller scales the situation becomes much more complicated anyway due to point sources, beam size and other non-linear effects.

Lensing induces non-Gaussianities in the lensed CMB sky, and also changes the power spectra of the perturbations. Lensing will start to have an observable effect on the power spectrum very shortly, and hence needs to be taken into account to obtain correct parameter constraints and error bars. For future observations, including the Planck[1] satellite and forthcoming ground based telescopes the effect is very important. In this paper we describe the simulation of lensed CMB maps (including the full non-Gaussian structure), but show that using an accurate calculation of the lensed CMB power spectra [4] a naive parameter estimation (neglecting non-Gaussianities) works rather well up to Planck sensitivities. Observations at higher sensitivities and resolutions will probably require a fuller analysis accounting for the full non-Gaussian distribution of the lensed sky, an important problem that we do not tackle here. Our simulation code can be used for testing future methods, and the simple power spectrum parameter estimation method can act as a useful baseline for future improvements.

II. WEAK LENSING OF THE CMB

The small scale CMB anisotropy is dominated by the emission from the last scattering surface at redshift $z \sim 1000$. Weak lensing of the CMB deflects photons coming from an original direction $\hat{\mathbf{n}}'$ on the last scattering surface to an observed direction $\hat{\mathbf{n}}$ on the sky today, so a lensed CMB field is given by

$$\widetilde{X}(\hat{\mathbf{n}}) = X(\hat{\mathbf{n}}') \tag{1}$$

where X is the unlensed field. The arcminute-scale displacement of the points is determined by the potential along the line of sight to the last scattering surface, conveniently encapsulated into an integrated lensing potential ψ [3]. The deflection vector is given by the gradient of the lensing potential $\boldsymbol{\nabla}\psi(\hat{\mathbf{n}})$, where $\boldsymbol{\nabla}$ is the covariant derivative on the sphere. The vector $\hat{\mathbf{n}}'$ is obtained from $\hat{\mathbf{n}}$ by moving its end on the surface of a unit sphere a distance $|\boldsymbol{\nabla}\psi(\hat{\mathbf{n}})|$ along a geodesic in the direction of $\boldsymbol{\nabla}\psi(\hat{\mathbf{n}})$ [3, 5]. This is sometimes written as $\hat{\mathbf{n}}' = \hat{\mathbf{n}} + \boldsymbol{\nabla}\psi(\hat{\mathbf{n}})$. To our level of approximation $|\boldsymbol{\nabla}\psi|$ is assumed to be constant between $\hat{\mathbf{n}}$ and $\hat{\mathbf{n}}'$, consistent with working out the lensing potential in the Born approximation (evaluating the potential along the undeflected path). Lensing deflections are a few arcminutes, but are coherent over degree scales, so this is a good approximation.

The power spectrum for the lensing potential $C_l^{\psi\psi}$ (and the correlation to the temperature $C_l^{\psi T}$) can be computed numerically in linear theory for a particular cosmological model using CAMB[2] [4, 6, 7]. From the power spectrum simulated realizations can be made assuming

*URL: http://cosmologist.info

[1] http://sci.esa.int/planck

[2] http://camb.info

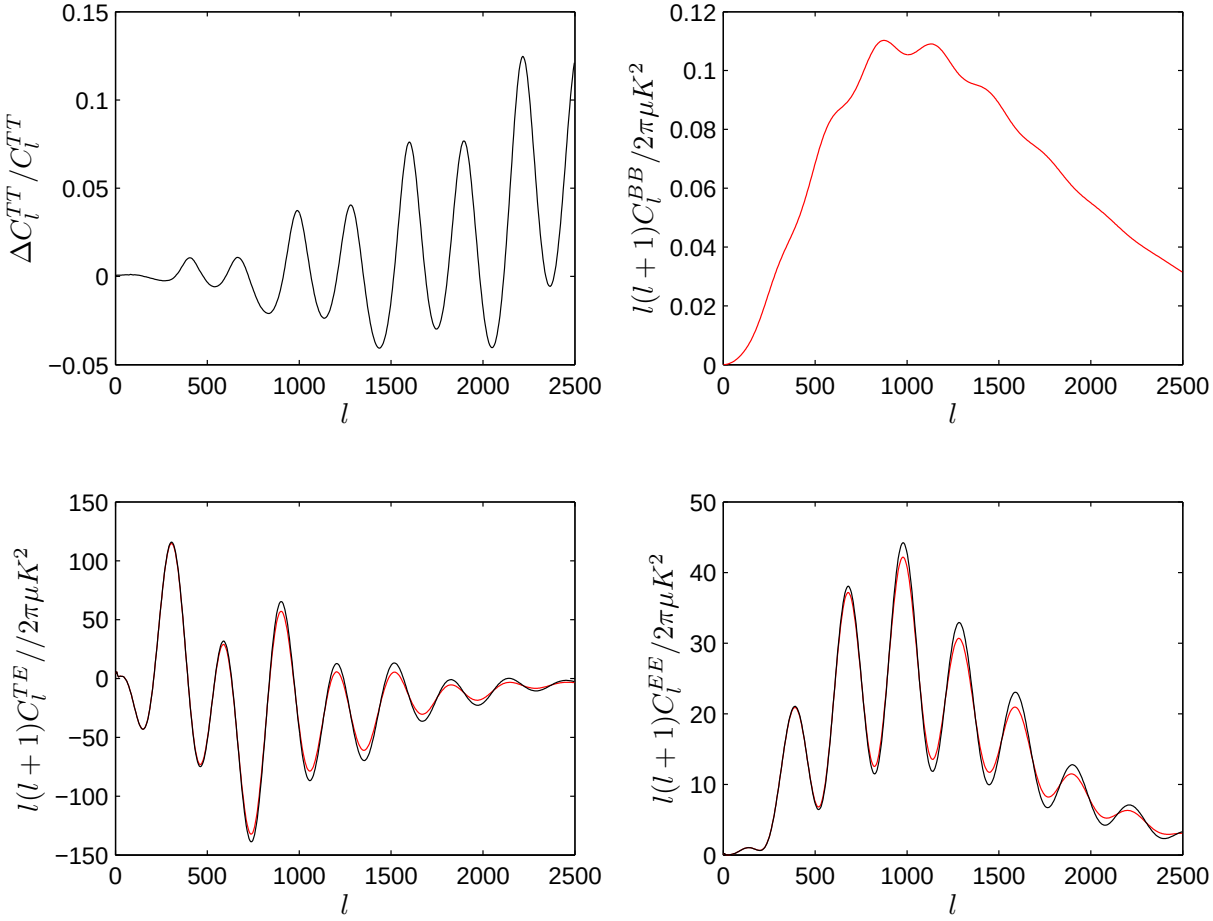

FIG. 1: The effect of lensing on the CMB power spectra. The top plots show the fractional change in the temperature spectrum C_l^{TT} and the lensing-induced B-polarization spectrum C_l^{BB}. The bottom plots show the lensed (grey/red, less peaked) and unlensed (black) T-E cross-correlation C_l^{TE} and E-polarization C_l^{EE} power spectra. All results are for the fiducial model given in the text, and the lensed B-mode power spectrum shown is not very accurate due to the neglect of non-linear evolution in the lensing potential.

Gaussianity of the primordial fields. Non-linear evolution of the potential changed the power spectrum $C_l^{\psi\psi}$ on small scales and also makes the ψ distribution somewhat non-Gaussian. The power spectrum $l^4 C_l^{\psi\psi}$ peaks at $l \sim 60$, however small scales where non-linear evolution is important at late times only contribute to ψ at $l \gg 60$. On these scales the contributions to the lensing potential come from a rather broad range of redshifts from $1 \lesssim z \lesssim 10$, so even at $l \sim 1000$ the late time non-linear evolution does not radically change $C_l^{\psi\psi}$. We therefore neglect the small effect of non-linear evolution here, though it can become important for high resolution polarization B-mode experiments. Using CAMB with HALOFIT [4, 8] the non-linear evolution of ψ can be estimated to change the lensed temperature power spectrum $\widetilde{C}_l^{TT}$ by about $\sim 0.2\%$ at $l \sim 2000$, though growing to 1% or more on smaller scales.

The significant effect of lensing on the CMB power spectra is shown in Fig. 1, where the lensed power spectra are computed accurately numerically using the correlation function method of Ref. [4]. The lensing smoothes out features in the temperature and polarization power spectra, changing the C_l peaks by up to 20% for the E-polarization on the scales of interest. Weak lensing does not change the total variance of the CMB anisotropies, with

$$\sum_l (2l+1)\widetilde{C}_l = \sum_l (2l+1)C_l, \qquad (2)$$

and similarly for the polarization [4]. This encapsulates that fact that the observation in any fixed direction is just a displaced view of the last scattering surface, and hence is Gaussian and with the same variance as if there were no lensing. Lensing induces a non-Gaussian spatial correlation structure to the lensed CMB fields but does not alter the variance at a point.

We now move on to discuss how to simulate full-sky lensed CMB maps.

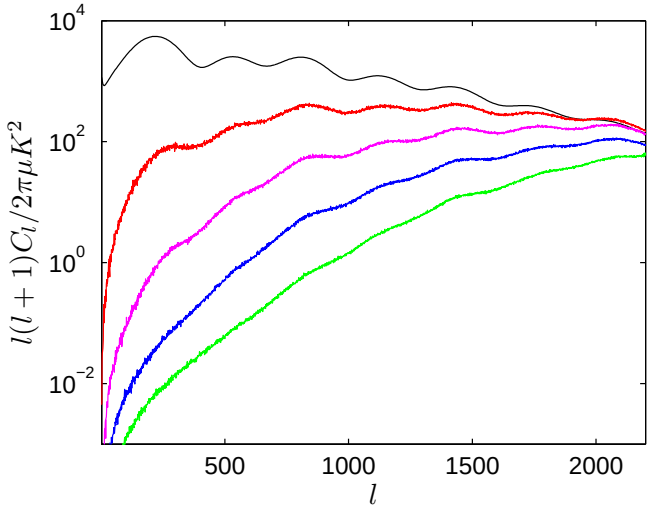

FIG. 2: Power spectra from realizations of the 1st, 2nd, 3rd and 4th order terms in the lensing potential series expansion of the lensed temperature $\widetilde{T}_{lm}$ compared to the full lensed $\widetilde{C}_l$ (top). The spectra may contain some pixelization error.

A. Series expansion

A common procedure for working with CMB weak lensing it to perform a series expansion in the deflection angle, so for the temperature

$$\begin{aligned}\widetilde{T}(\hat{\mathbf{n}}) &= T(\hat{\mathbf{n}}+\boldsymbol{\nabla}\psi) \\ &= T+\boldsymbol{\nabla}_a T\boldsymbol{\nabla}^a\psi+\tfrac{1}{2}\boldsymbol{\nabla}^a\psi\boldsymbol{\nabla}^b\psi\boldsymbol{\nabla}_a\boldsymbol{\nabla}_b T+\ldots \quad (3)\end{aligned}$$

and the last line is evaluated at $\hat{\mathbf{n}}$. This expansion will be valid for scales much smaller than the deflection angle. On smaller scales the unlensed fields are deflected by a distance comparable to their wavelength, and the change in phase is not a small perturbation. Even on scales where the expansion is valid, the expansion only converges relatively slowly. The power spectra of the individual terms in the expansion are shown in Fig. 2 up to fourth order. This clearly shows that a low order series expansion cannot be used for accurate map simulation due to the large variance of higher order terms on small scales. The lowest order series expansion may be a useful approximation for some applications (e.g. Ref. [9]), but can easily give results which are sufficiently inaccurate to be problematic. For example it was shown in Ref. [4] that the lensed CMB power spectrum computed from the lowest order series expansion of Ref. [7] gives lensing corrections which are incorrect by an order unity factor on small scales.

B. Re-mapping points

Since the series expansion is not accurate enough for making simulated maps at high resolution, the best way to proceed is to re-map the anisotropy field by the deflection vector $\boldsymbol{\nabla}\psi$ as a function of position. The deflection field is a vector field on the sphere, and can be easily simulated using spin-1 spherical harmonics (or, equivalently, vector spherical harmonics). Full details of the simulation process for the temperature and polarization on the sphere are given in the Appendix, along with details of a fast multi-processor implementation using a modified version of HEALPIX[3] 1.2 [10].

It is important to remember that the lensing deflects the physical anisotropy field, not the field after beam and pixel convolution. For this reason sky simulation and re-mapping have to be performed at high resolution regardless of how broad the observational beam is. The lensed field can then be convolved as required.

III. PARAMETER ESTIMATION

With future precision observations we would ideally like to extract information about both the unlensed CMB (which has simple Gaussian statistical properties), and about the lensing potential (which contains additional information). Gravitational lensing of scalar E-mode polarization may ultimately dominate the B-mode polarization signal from primordial gravitational waves, so an accurate treatment of the lensing will become essential [11, 12, 13]. In the more immediate future gravitational lensing will have a significant effect on the statistics of the observed CMB, and must be accounted for to obtain reliable parameter estimates. The non-Gaussianities induced by lensing must also be accounted for when attempting to assess the degree of primordial non-Gaussianity (an important probe of early universe physics).

The amount of information that can be learned about the lensing potential from the non-Gaussianities depends on the noise level [9, 14, 15]. For observations up to Planck sensitivity any reconstructed map of the lensing potential would be completely noise dominated, and the extra information this contains is rather limited. However the effect of lensing on the power spectrum is significant and certainly cannot be ignored. In this paper we do not address the problem of handling the full likelihood function, but merely show that using the lensed power spectrum and approximating the lensed field as Gaussian is sufficient to obtain good parameter estimates at Planck sensitivity. More general methods for handling the lensing likelihood function will be discussed in a future paper if they can be made to work.

[3] http://www.eso.org/science/healpix/

A. Gaussian C_l Likelihood function

Spherical harmonic coefficients of a Gaussian temperature and polarization field on the sky can be used to define estimators of the covariance

$$\hat{C}_l^{WX} = \frac{1}{2l+1}\sum_m W_{lm}^* X_{lm} \tag{4}$$

where W_{lm} and X_{lm} are spherical harmonic coefficients of the temperature T, E-polarization E or B-polarization B. These are unbiased in that averaged over realizations $\langle \hat{C}_l^{WX}\rangle = C_l^{WX}$.

The assumed Gaussianity of T_{lm}, E_{lm} and B_{lm} gives the following full-sky likelihood function:

$$-2\log P(\hat{\boldsymbol{C}}_l|\boldsymbol{C}_l) = (2l+1)\left\{\mathrm{Tr}\left[\hat{\boldsymbol{C}}_l\boldsymbol{C}_l^{-1}\right] + \log|\boldsymbol{C}_l|\right\} \tag{5}$$

(to within an irrelevant constant) where

$$\boldsymbol{C}_l = \begin{pmatrix} C_l^{TT} & C_l^{TE} & 0 \\ C_l^{TE} & C_l^{EE} & 0 \\ 0 & 0 & C_l^{BB}\end{pmatrix} \tag{6}$$

and $\hat{\boldsymbol{C}}$ is the corresponding matrix of estimators. In the presence of instrumental noise the C_l and $\hat{C}_l$ should include the noise variance. We have assumed a statistically parity invariant ensemble so that $C^{BT} = C^{BE} = 0$.

For the lensed sky this is not the correct relation because the lensed sky is not Gaussian if the lensing potential is not fixed. However as discussed below, replacing C_l by $\widetilde{C}_l$ does not give significantly biased results for Planck, and is by far the simplest method of accounting for CMB lensing in a parameter analysis: i.e. just pretend the lensed field is Gaussian, and use the theoretical lensed CMB power spectrum. However at high resolutions and sensitivity this will not be correct, as the statistics of the $\widetilde{C}_l$ on small scales are governed by the same small number of large scale lensing modes [16]. The naive approach could be improved by (for example) calculating an empirical $\hat{\widetilde{C}}_l$ covariance from simulations and using this to make some partial correction for non-Gaussianity induced variations to the posterior distribution. Non-Gaussian corrections should be small at Planck sensitivity, however the non-Gaussian corrections to futuristic signal dominated lensing B-mode observations can be important [17]. We have checked that the lensed estimator $\hat{\widetilde{C}}_l{}^{TT}$ and $\hat{\widetilde{C}}_l{}^{EE}$ variances from simulations agree with the Gaussian results to within a few percent at $l < 2000$.

B. Sampling from the posterior

We assume a simple flat adiabatic ΛCDM cosmological model with the following parameters to be determined from the data: primordial curvature perturbation power spectrum with spectral index n_s and amplitude A_s (at wavenumber $0.05\mathrm{Mpc}^{-1}$), baryon density $\Omega_b h^2$, cold dark matter density $\Omega_c h^2$, optical depth τ (reionization assumed sharp), and Hubble parameter today $H_0 = 100h$ km s^{-1}Mpc^{-1}. We approximate the neutrinos as massless and assume standard general relativity.

We use the CosmoMC[4] [18] Markov Chain Monte Carlo (MCMC) package to sample from the posterior distribution of the parameters given the observed (simulated) data. To make the posterior parameter distributions more Gaussian we use θ_r as a base parameter (with flat prior) instead of the Hubble parameter H_0. The derived parameter θ_r is defined as the (approximate) ratio of the sound horizon at last scattering to the angular diameter distance [19], a non-linear function of the other parameters that is very well constrained by the position of the acoustic peaks. We also transform to the amplitude parameter $\ln A_s$ (with a flat prior) that is constrained very well in a linear combination with τ because $A_s e^{-2\tau}$ determines the small scale amplitude of the C_l. We then use the covariance matrix to transform to an uncorrelated set of parameters that the MCMC proposal density can use to explore the posterior distribution efficiently. Given an approximate covariance matrix from previous runs chains converge in a few hours. For the first run the covariance can be learned dynamically as the chain evolves (discarding samples from the evolving part of the chain), or one can use the Hessian at the best fit point as a useful starting approximation.

Since the computation time for generating the lensed $\widetilde{C}_l$ is dominated by the time to compute the transfer functions for the unlensed C_l and the lensing potential, parameters like $\ln A_s$ and n_s remain 'fast' parameters [18], in that changing them is quick as long as the other parameters remain fixed. Thus methods to efficiently exploit the difference between 'fast' and 'slow' parameters [20] can still be used to speed up MCMC runs even when CMB lensing is included.

We neglect $\widetilde{C}_l^{BB}$ since it is noise dominated even at Planck sensitivity and has almost no effect on parameter constraints, though including it in the Gaussian approximation is trivial if desired.

C. WMAP data

The importance of the lensing effect depends on the amplitude of the potential fluctuations along the line of sight. Larger amplitudes cause more lensing. By itself the first year WMAP data [21] constrains the amplitude rather poorly due to a degeneracy with the optical depth and the absence of C_l^{EE} data to give a good upper limit. Large values of the optical depth $\tau \gtrsim 0.5$ allowed by the data correspond to models with large amplitudes in

[4] `http://cosmologist.info/cosmomc`

which the lensing effect is quite significant, so strictly speaking the lensing should be accounted for in analyses of the first year WMAP data alone [22, 23]. In a totally free WMAP first year parameter analysis the lensing already has a noticeable effect. However the region of parameter space with $\tau \gtrsim 0.5$ is almost certainly disallowed by numerous other sources of data and so should not be taken too seriously.

The second year data WMAP should include C_l^{EE} and a better measurement of the third peak, restricting to a much smaller range of τ and $\Omega_b h^2$, and so making the effect of lensing on the tails of the distribution less significant. As a simple toy model consider a full sky observation with isotropic Gaussian noise with variance $N_l^{TT} = N_l^{EE}/4 = N_l^{BB}/4 = 0.03\mu K^2$ and a symmetric Gaussian beam of 13 arcminutes full-width half-maximum and neglect foregrounds. We find that neglecting the lensing effect on the power spectrum leads to a posterior mean of $\Omega_b h^2$ about half a standard deviation lower than it should be for a typical realization. This is easily understood: lensing smoothes out the third peak in the temperature C_l, making it appear lower relative to the first and second peaks by a couple of percent. A similar effect can be produced without lensing by lowering the baryon density, so unless lensing is modelled consistently there is a danger of confusion.

Future ground or balloon results that can resolve the higher temperature C_l peaks will also be sensitive to lensing as the corrections become larger than 5% at $l \gtrsim 1500$. We conclude that very soon it will be important to include lensing in parameter analyses to obtain accurate results from CMB observations.

D. Planck-like simulation

As an example of a future observation where modelling the lensing will be crucial we now consider a simple full-sky CMB observation simulation at Planck-like sensitivity. We compute theoretical CMB power spectra for the lensing potential and unlensed fields for a simple fiducial model with $n_s = 0.99$, $A_s = 2.5 \times 10^{-9}$, $\Omega_b h^2 = 0.22$, $\Omega_c h^2 = 0.12$, $\tau = 0.15$, $h = 0.72$, and simulate full sky lensed maps with $12 \times 2048^2 \sim 5 \times 10^7$ pixels (generated by remapping a $12 \times 8192^2 \sim 8 \times 10^8$ pixel unlensed sky as described in the appendix). As a simple toy model of an observation at optimistic Planck-like sensitivity we assume isotropic Gaussian noise with variance $N_l^{TT} = N_l^{EE}/4 = N_l^{BB}/4 = 2 \times 10^{-4}\mu K^2$ and a symmetric Gaussian beam of 7 arcminutes full-width half-maximum and neglect foregrounds.

We use simulated lensed $\widetilde{C}_l$ on scales with $l \leq 2000$. As discussed above we neglect non-Gaussianity of the lensed sky, but account for the lensing by using accurate theoretical lensed CMB power spectra [4]. Obtaining parameter estimates at Planck sensitivity is then no more difficult than with WMAP.

Fig. 3 shows the posterior parameter constraints from this particular realization. The constraints from the lensed sky analysed using the lensed $\widetilde{C}_l$ are in quite good agreement with those one would have obtained if the unlensed sky were observable and was analysed with the unlensed C_l. So current tools appear to be sufficient to extract parameter constraints reliably at up to Planck-like sensitivity. Of course it is essential to model the lensing consistently: Fig. 3 also demonstrates the wrong parameter constraints that are obtained if the lensed sky is incorrectly analysed using the unlensed C_l.

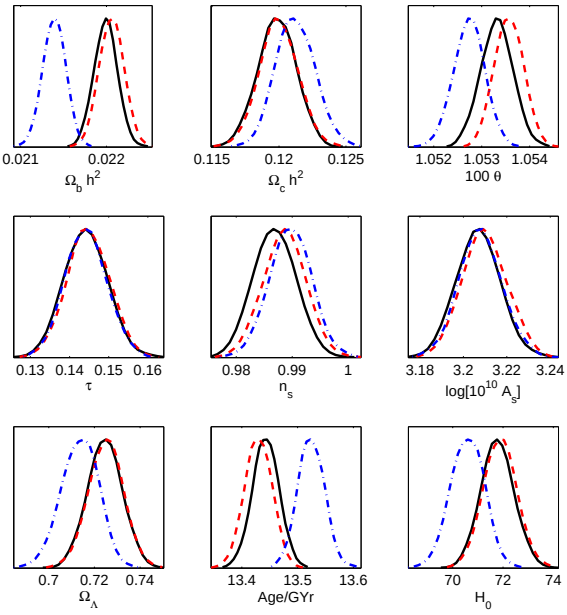

FIG. 3: Parameter constraints from a simple Planck-like simulation. Solid lines analyse the lensed sky assuming Gaussianity with the lensed CMB power spectra, dashed lines are for an unlensed sky analysed with the unlensed power spectra, dash-dotted lines show the (inconsistent) result from analysing the lensed sky using the unlensed power spectra. The bottom row shows the dark energy density, age and Hubble constant derived parameters.

We appear to agree with Ref. [24] that at this sensitivity lensing does little to change the effectiveness of parameter estimation, with the recovered error bars being of approximately the same width as when analysing an unlensed sky. However we have not tested the accuracy of the recovered error bars, merely showing that the posterior peaks are at about the correct parameter values consistent with the error bar. The result of Ref. [16] would suggest that the error bars are likely to be correct to $\lesssim 5\%$, but for a precision estimate of the error bars a fuller analysis would be required. Increasing the error bars should make the lensed and unlensed analysis results even more consistent, though there is no reason in general why the posterior means should be identical in any given realization.

The effect of non-linear evolution of the lensing potential can be safely neglected at Planck sensitivity. How-

ever we have made one important linear-theory assumption, which is that the unlensed CMB power spectra can be worked out accurately. Given a standard cosmology and an ionization history this is straightforward to do, however uncertainties in the complicated details of recombination may mean that the ionization history is not very well known[5] [25, 26, 27]. Future CMB parameter analyses, including Planck, may require a more accurate calculation of recombination to obtain unlensed (and hence lensed) power spectra accurately enough. In addition, other second order signatures such as from the kinetic Sunyaev-Zel'dovich effect will probably have a small but important contribution that also needs to be accounted for [28, 29, 30, 31].

IV. CONCLUSIONS

We have described a lensed CMB sky simulation code that can be used for testing future analysis methods. As a simple benchmark we have demonstrated parameter estimation at up to Planck sensitivity by using the lensed CMB power spectrum on scales $l \leq 2000$. This recovers the correct parameters, though the error bars may be less accurate. More sensitive observations will require a fuller analysis of the non-Gaussian likelihood function. Higher resolution observations may require a much more complicated numerical analysis of non-linear evolution of the potential and other non-linear effects [29, 32].

Incidentally we have also demonstrated that current parameter estimation methodology is sufficiently accurate for parameter estimation with Planck (under the assumption of linearity and assuming that the ionization history is known well enough). The simple lensing analysis considered here only causes an order unity increase in computing time compared to an unlensed analysis, so there is no problem accounting for the lensing effect in parameter analyses. Future CMB parameter studies must account for CMB lensing to obtain correct results. The lensed CMB simulation code is publicly available[6], as is the parameter estimation code including support for lensed CMB power spectra[7] [18].

Acknowledgments

I thank Matias Zaldarriaga, Anthony Challinor, Carlo Contaldi, Ben Wandelt, Mike Nolta and Sarah Smith for discussion and comments. Some of the results in this paper have been derived using a modified version of the HEALPIX [10] package. The beowulf computer used for some of this analysis was funded by the Canada Foundation for Innovation and the Ontario Innovation Trust.

APPENDIX A: HARMONICS AND MAP MAKING

We use spherical polar coordinates, with orthonormal basis vectors at any point on the sphere given by $\mathbf{e}_\theta$ and $\mathbf{e}_\phi$. A complex spin-s quantity can be defined as the components of a rank-$|s|$ tensor in the basis $\mathbf{e}_\pm \equiv \mathbf{e}_\theta \pm i\mathbf{e}_\phi$. For example the spin two polarization [33] is given in terms of the polarization tensor P_{ab} by ${}_2P = \mathbf{e}_+^a \mathbf{e}_+^b P_{ab}$. Further definitions and derivations using our notation can be found in the appendix of Ref. [34].

A spin s quantity ${}_s\eta$ has harmonic components ${}_sa_{lm}$ given by

$$ {}_s\eta = \sum_{lm} {}_sa_{lm}\, {}_sY_{lm} \tag{A1}$$

where $l \geq |m|, l \geq |s|$. The spin harmonics are defined by

$$ {}_sY_{lm} \equiv {}_s\lambda_{lm} e^{im\phi} \equiv \sqrt{\frac{(l-|s|)!}{(l+|s|)!}}\ \eth^s Y_{lm}, \tag{A2}$$

and $\eth^{-|s|} \equiv (-1)^s \overline{\eth}^{|s|}$, Y_{lm} is a standard spin zero harmonic and $\eth$ is the spin raising operator (see Ref. [34]). The harmonics have the symmetries

$$\begin{aligned} {}_sY^*_{lm} &= (-1)^{s+m}\, {}_{-s}Y_{l(-m)} \\ {}_sY_{lm}(\pi-\theta,\phi) &= (-1)^{l+m}\, {}_{-s}Y_{lm}(\theta,\phi). \end{aligned} \tag{A3}$$

We define gradient (E) and curl (B) harmonic components of a tensor field with the general definition

$$\begin{aligned} {}_{|s|}E_{lm} &\equiv (-1)^H \tfrac{1}{2}({}_{|s|}a_{lm} + (-1)^s\, {}_{-|s|}a_{lm}) \\ i\, {}_{|s|}B_{lm} &\equiv (-1)^H \tfrac{1}{2}({}_{|s|}a_{lm} - (-1)^s\, {}_{-|s|}a_{lm}) \end{aligned} \tag{A4}$$

where ${}_{|s|}E^*_{lm} = (-1)^m\, {}_{|s|}E_{l(-m)}$, ${}_{|s|}B^*_{lm} = (-1)^m\, {}_{|s|}B_{l(-m)}$ and $(-1)^H$ is a sign convention. This definition ensures that gradient fields are always pure E. In general a complex spin field ${}_s\eta$, with ${}_{|s|}\eta^* = {}_{-|s|}\eta$, can be expanded as

$$\begin{aligned} {}_{|s|}\eta &= (-1)^H \sum ({}_{|s|}E_{lm} + i\, {}_{|s|}B_{lm})\, {}_{|s|}Y_{lm} \\ {}_{-|s|}\eta &= (-1)^{H+s} \sum_{lm} ({}_{|s|}E_{lm} - i\, {}_{|s|}B_{lm})\, {}_{-|s|}Y_{lm}. \end{aligned} \tag{A5}$$

For polarization ${}_2P = \mathbf{e}_+^a \mathbf{e}_+^b P_{ab} = Q + iU$ where Q and U are the Stokes parameters measured in the $(\mathbf{e}_\theta, \mathbf{e}_\phi)$ basis. Where Q and U are instead measured with respect to a left handed set $(\mathbf{e}_\theta, -\mathbf{e}_\phi)$ as in Refs [34, 35] we have ${}_2P = Q - iU$ (this basis defines a right handed set about the incoming photon direction). In HEALPIX[8]

[5] http://cosmocoffee.info/viewtopic.php?t=174
[6] http://cosmologist.info/lenspix
[7] http://cosmologist.info/cosmomc
[8] http://www.eso.org/science/healpix/

1.2 conventions ${}_2P = Q + iU$, and $H = 1$, whereas in the conventions of Refs [34, 35] $H = 0$. On large scales $C_l^{TE} = (-1)^{H+1}|C_l^{TE}|$, and the $H = 1$ convention is used for the output from CAMB and CMBFAST so the large scale correlation is positive. To be consistent with these codes we use the $H = 1$ convention for numerical work.

Harmonic transforms

For CMB lensing we need to generate maps of the gradient of the potential, a vector field on the sphere. Various first order results (for example the efficient quadratic estimators for the potential [9]) also require transforming to and from spin one and spin three fields.

To compute the spin s harmonics one can either iterate the spin s recursion relation (see Ref. [34]), or relate the harmonics to the spin zero harmonics. Here we give the results for the latter approach, which is followed by HEALPIX [10]. Transforms of spin zero fields and spin two fields (for polarization) are standard [35] and included in HEALPIX 1.2.

Defining

$$ {}_sF^{\pm}_{lm} \equiv \sqrt{\frac{(l-s)!}{(l+s)!}}\,{}_sW^{\pm}_{lm} \equiv \tfrac{1}{2}({}_s\lambda_{lm} \pm (-1)^s\,{}_{-s}\lambda_{lm}) = \tfrac{1}{2}({}_s\lambda_{lm} \pm (-1)^m\,{}_s\lambda_{l(-m)}) \quad \text{(A6)} $$

and using some results from Ref. [34] we have

$$ \begin{aligned}
{}_1W^+_{lm} &= \frac{1}{\sin\theta}\left[\alpha_{lm}\lambda_{(l-1)m} - l\cos\theta\lambda_{lm}\right] \\
{}_1W^-_{lm} &= \frac{m}{\sin\theta}\lambda_{lm} \\
{}_2W^+_{lm} &= \left[\frac{2(m^2-l)}{\sin^2\theta} - l(l-1)\right]\lambda_{lm} + \frac{2\cos\theta}{\sin^2\theta}\alpha_{lm}\lambda_{(l-1)m} \\
{}_2W^-_{lm} &= \left[\alpha_{lm}\lambda_{(l-1)m} - (l-1)\cos\theta\lambda_{lm}\right]\frac{2m}{\sin^2\theta} \\
{}_3W^+_{lm} &= \left[\left(l(l-1)(l-2) - 4\frac{2l+m^2(l-3)}{\sin^2\theta}\right)\cos\theta\lambda_{lm}\right. \\
&\quad \left. -\alpha_{lm}\left(l(l+1) + 6 - \frac{4(2+m^2)}{\sin^2\theta}\right)\lambda_{(l-1)m}\right]\frac{1}{\sin\theta} \\
{}_3W^-_{lm} &= \left[\left(4\frac{m^2-3l+2}{\sin^2\theta} - 3(l-1)(l-2)\right)\lambda_{lm}\right. \\
&\quad \left. +12\alpha_{lm}\frac{\cos\theta\lambda_{(l-1)m}}{\sin^2\theta}\right]\frac{m}{\sin\theta}, \quad \text{(A7)}
\end{aligned} $$

where

$$ \alpha_{lm} \equiv \sqrt{\frac{(2l+1)(l^2-m^2)}{2l-1}}. \quad \text{(A8)} $$

Similar results can be derived for higher spins if desired. If several different spin transforms are being done at the same time one can also use a relation like

$$ l\sin\theta\,{}_{s+1}W^{\pm}_{lm} = (l-s)\left(m\,{}_sW^{\mp}_{lm} - l\cos\theta\,{}_sW^{\pm}_{lm}\right) + (l+s)\alpha_{lm}\,{}_sW^{\pm}_{(l-1)m}. \quad \text{(A9)} $$

To transform to and from a map of a spin field ${}_{|s|}\eta = R + iI$ (where R and I are real) we use the symmetry in θ so that

$$ \begin{aligned}
R(\theta,\phi) &= (-1)^H\sum_{lm}\left({}_sF^+_{lm|s|}E_{lm} + i_sF^-_{lm|s|}B_{lm}\right)e^{im\phi} \\
I(\theta,\phi) &= (-1)^H\sum_{lm}\left({}_sF^+_{lm|s|}B_{lm} - i_sF^-_{lm|s|}E_{lm}\right)e^{im\phi} \\
R(\pi-\theta,\phi) &= (-1)^{H+l+m+s}\times \\
&\quad \sum_{lm}\left({}_sF^+_{lm|s|}E_{lm} - i_sF^-_{lm|s|}B_{lm}\right)e^{im\phi} \\
I(\pi-\theta,\phi) &= (-1)^{H+l+m+s}\times \\
&\quad \sum_{lm}\left({}_sF^+_{lm|s|}B_{lm} + i_sF^-_{lm|s|}E_{lm}\right)e^{im\phi}. \quad \text{(A10)}
\end{aligned} $$

For pixelizations where pixels are on lines of constant latitude (such as HEALPIX, IGLOO [36] and GLESP [37]) the ϕ transform can be performed rapidly using FFTs. The remaining computational cost is easily parallelized by sending pixels at different θ to separate processors, and $l_{\rm max} \approx 2000$, $n_{\rm pix} \approx 10^7$ maps can be generated in a few seconds over about fifty modern processors.

The deflection field

The gradient of the potential (a scalar), $\boldsymbol{\nabla}\psi$, has harmonic components

$$ {}_1E_{lm} = (-1)^{H+1}\sqrt{l(l+1)}\psi_{lm} \qquad {}_1B_{lm} = 0 \quad \text{(A11)} $$

which follows from

$$ \begin{aligned}
{}_1a_{lm} &= \int d\Omega(\mathbf{e}_+\cdot\boldsymbol{\nabla}\psi)\,{}_1Y^*_{lm} \\
&= -\int d\Omega\,\eth\,\psi\,{}_1Y^*_{lm} = -\sqrt{l(l+1)}\psi_{lm} \quad \text{(A12)}
\end{aligned} $$

$$ \begin{aligned}
{}_{-1}a_{lm} &= \int d\Omega(\mathbf{e}_-\cdot\boldsymbol{\nabla}\psi)\,{}_{-1}Y^*_{lm} \\
&= -\int d\Omega\,\overline{\eth}\,\psi\,{}_{-1}Y^*_{lm} = \sqrt{l(l+1)}\psi_{lm} \quad \text{(A13)}
\end{aligned} $$

where $\mathbf{e}_\pm\cdot\boldsymbol{\nabla}\psi$ are the spin ±1 components of $\boldsymbol{\nabla}\psi$. Thus maps of the gradient field can easily be constructed from the harmonic components of ψ using results from the last section. For spin one field ${}_1\eta = R + iI = \mathbf{e}_+\cdot\mathbf{X}$, we see that R and I are simply the $\mathbf{e}_\theta$ and $\mathbf{e}_\phi$ components of the vector field $\mathbf{X}$.

Accurate lensed map making

To make accurate lensed temperature maps on the full sky we use

$$\widetilde{T}(\hat{\mathbf{n}}) = T(\hat{\mathbf{n}}') = \sum_{lm} T_{lm} Y_{lm}(\hat{\mathbf{n}}'). \tag{A14}$$

Given a set of harmonic coefficients, constructing the lensed map is straightforward. Using identities for spherical triangles, the lensed temperature at a position (θ, ϕ) is given by the unlensed temperature at position $(\theta', \phi+\Delta\phi)$ where

$$\cos\theta' = \cos d\cos\theta - \sin d\sin\theta\cos\alpha \tag{A15}$$

$$\sin\Delta\phi = \frac{\sin\alpha\sin d}{\sin\theta'} \tag{A16}$$

and the deflection vector is $\mathbf{d} \equiv \boldsymbol{\nabla}\psi = d_\theta \mathbf{e}_\theta + d_\phi \mathbf{e}_\phi = d\cos\alpha\mathbf{e}_\theta + d\sin\alpha\mathbf{e}_\phi$. Except near the coordinate singularities the obvious Euclidean results are a rather accurate approximation.

For the polarization the points move the same way, with the polarization maintaining the same orientation relative to the deflection vector at the two points (neglecting field rotation, see Ref. [11]). However for the components of a spin field we have to be careful to account for the different direction of the coordinate vectors at the two points [5]. This requires rotating the components of the spherical polar coordinates by γ, the difference between the angle made by $\mathbf{e}_\theta$ and the connecting geodesic at the two points. If $\alpha' \equiv \alpha - \gamma$ we have

$$\tan(\alpha') = \frac{d_\phi}{d\sin d\cot\theta + d_\theta\cos d} \tag{A17}$$

and after weak lensing a spin s field ${}_s\eta$ becomes

$${}_s\widetilde{\eta}(\hat{\mathbf{n}}) = e^{is\gamma}{}_s\eta(\hat{\mathbf{n}}'). \tag{A18}$$

For the spin two polarization field we can avoid inverse trigonometric functions by using

$$e^{2i\gamma} = \frac{2(d_\theta + d_\phi A)^2}{d^2(1+A^2)} - 1 + \frac{2i(d_\theta + d_\phi A)(d_\phi - d_\theta A)}{d^2(1+A^2)} \tag{A19}$$

where $A \equiv \tan\alpha'$. Except near the poles this is very close to unity.

Note that to get an accurate simulation of the lensing B modes at $l \gtrsim 1000$ it is necessary to include relatively high l (more than just $l_{\rm max}+500$ that is accurate for the other spectra). The B lensing signal only becomes useful for parameter estimation after Planck.

To generate the lensed field we can not use FFTs because even if $\hat{\mathbf{n}}$ is sampled equally in ϕ on rings of constant θ, the original positions $\hat{\mathbf{n}}'$ will not. An $l_{\rm max} \approx 2000$, $n_{\rm pix} \approx 10^7$ polarized map can be made in about 2000 CPU hours, with good trivial parallelization. In practice a much faster way to make maps to good accuracy is to generate an unlensed sky at higher resolution and just re-map the points. We find that using $12\times 8192^2 \sim 10^9$ pixels for the high resolution map is sufficient to get a 12×1024^2-pixel lensed sky with $\widetilde{C}_l$ accurate to 0.5% at $l \lesssim 2000$. The fractional accuracy on the polarization is similar, except for the B modes induced by lensing which are a more sensitive: $12\times16384^2 \sim 3\times10^9$ pixels are required to get $\widetilde{C}_l^{BB}$ accurate at percent level for $l \lesssim 2000$. The B-spectrum is quite sensitive to $l_{\rm max}$ and the non-linear power spectrum, so this accuracy level is only notional. Note that choosing $l_{\rm max}$ too low generally underestimates the B-mode power, whereas pixelization errors lead to an overestimation, so it is possible to get spuriously accurate power spectra without actually having computed the lensed sky accurately.

The approximate point-remapping method is also easily parallelized, and by only generating sections of the high-resolution map on each cluster node, the total memory requirement per node can remain below one gigabyte as long as enough nodes are available. The scaling is then approximately the same as making maps at the higher resolution, and even multi-billion pixel remappings can be done in under an hour. For Planck resolution observations, lensed simulations can be done in a few minutes with enough processors. For applications where this speed is an issue some kind of faster interpolation scheme might be useful (avoiding the need to generate a fine map at much higher resolution than the base pixelization).

The parallelized code incidentally also provides a fast method for performing spherical harmonic transforms (without lensing) on computer clusters using MPI.

[1] U. Seljak, Astrophys. J. **463**, 1 (1996), astro-ph/9505109.
[2] M. Zaldarriaga and U. Seljak, Phys. Rev. **D58**, 023003 (1998), astro-ph/9803150.
[3] W. Hu, Phys. Rev. **D62**, 043007 (2000), astro-ph/0001303.
[4] A. Challinor and A. Lewis (2005), astro-ph/0502425.
[5] A. Challinor and G. Chon, Phys. Rev. **D66**, 127301 (2002), astro-ph/0301064.
[6] A. Lewis, A. Challinor, and A. Lasenby, Astrophys. J. **538**, 473 (2000), astro-ph/9911177.
[7] W. Hu, Phys. Rev. **D62**, 043007 (2000), astro-ph/0001303.
[8] R. E. Smith et al. (The Virgo Consortium), Mon. Not. Roy. Astron. Soc. **341**, 1311 (2003), astro-ph/0207664.
[9] T. Okamoto and W. Hu, Phys. Rev. **D67**, 083002 (2003), astro-ph/0301031.
[10] K. M. Gorski et al., Astrophys. J. **622**, 759 (2005), astro-ph/0409513.
[11] C. M. Hirata and U. Seljak, Phys. Rev. **D68**, 083002 (2003), astro-ph/0306354.

[12] L. Knox and Y.-S. Song, Phys. Rev. Lett. **89**, 011303 (2002), astro-ph/0202286.
[13] M. Kesden, A. Cooray, and M. Kamionkowski, Phys. Rev. Lett. **89**, 011304 (2002), astro-ph/0202434.
[14] M. Kesden, A. Cooray, and M. Kamionkowski, Phys. Rev. **D67**, 123507 (2003), astro-ph/0302536.
[15] C. M. Hirata and U. Seljak, Phys. Rev. **D67**, 043001 (2003), astro-ph/0209489.
[16] W. Hu, Phys. Rev. **D64**, 083005 (2001), astro-ph/0105117.
[17] K. M. Smith, W. Hu, and M. Kaplinghat, Phys. Rev. **D70**, 043002 (2004), astro-ph/0402442.
[18] A. Lewis and S. Bridle, Phys. Rev. **D66**, 103511 (2002), astro-ph/0205436.
[19] A. Kosowsky, M. Milosavljevic, and R. Jimenez, Pys. Rev. D **66**, 063007 (2002), astro-ph/0206014.
[20] A. Lewis and S. Bridle, `http://cosmologist.info/notes/cosmomc.ps.gz`.
[21] D. N. Spergel et al., Astrophys. J. Suppl. **148**, 175 (2003), astro-ph/0302209.
[22] S. L. Bridle, A. M. Lewis, J. Weller, and G. Efstathiou, MNRAS **342**, L72 (2003), astro-ph/0302306.
[23] M. Tegmark et al. (SDSS), Phys. Rev. **D69**, 103501 (2004), astro-ph/0310723.
[24] M. Zaldarriaga, D. N. Spergel, and U. Seljak, Astrophys. J. **488**, 1 (1997), astro-ph/9702157.
[25] S. Seager, D. D. Sasselov, and D. Scott, Astrophys. J. Suppl. **128**, 407 (2000), astro-ph/9912182.
[26] P. K. Leung, C.-W. Chan, and M.-C. Chu, Mon. Not. Roy. Astron. Soc. **349**, 632 (2004), astro-ph/0309374.
[27] V. K. Dubrovich and S. I. Grachev (2005), astro-ph/0501672.
[28] L. Knox, R. Scoccimarro, and S. Dodelson, Phys. Rev. Lett. **81**, 2004 (1998), astro-ph/9805012.
[29] A. Amblard, C. Vale, and M. J. White, New Astron. **9**, 687 (2004), astro-ph/0403075.
[30] M. G. Santos, A. Cooray, Z. Haiman, L. Knox, and C.-P. Ma, Astrophys. J. **598**, 756 (2003), astro-ph/0305471.
[31] O. Zahn, M. Zaldarriaga, L. Hernquist, and M. McQuinn (2005), astro-ph/0503166.
[32] M. J. White and C. Vale, Astropart. Phys. **22**, 19 (2004), astro-ph/0312133.
[33] M. Zaldarriaga and U. Seljak, Phys. Rev. **D55**, 1830 (1997), astro-ph/9609170.
[34] A. Lewis, A. Challinor, and N. Turok, Phys. Rev. **D65**, 023505 (2002), astro-ph/0106536.
[35] M. Kamionkowski, A. Kosowsky, and A. Stebbins, Phys. Rev. **D55**, 7368 (1997), astro-ph/9611125.
[36] R. Crittenden and N. Turok (1998), astro-ph/9806374.
[37] A. G. Doroshkevich et al. (2005), astro-ph/0501494.